A parallelizable sampling method for parameter inference of large biochemical reaction models

Jan Mikelson and Mustafa Khammash
(Dated: June 28, 2016)

The development of mechanistic models of biological systems is a central part of Systems Biology. One major task in developing these models is the inference of the correct model parameters. Due to the size of most realistic models and their possibly complex dynamical behaviour one must usually rely on sample based methods. In this paper we present a novel algorithm that reliably estimates model parameters for deterministic as well as stochastic models from trajectory data. Our algorithm samples iteratively independent particles from the level sets of the likelihood and recovers the posterior from these level sets. The presented approach is easily parallelizable and, by utilizing density estimation through Dirichlet Process Gaussian Mixture Models, can deal with high dimensional parameter spaces. We illustrate that our algorithm is applicable to large, realistic deterministic and stochastic models and succeeds in inferring the correct posterior from a given number of observed trajectories. This algorithm presents a novel, computationally feasible approach to identify parameters of large biochemical reaction models based on sample path data.

arXiv:1606.08281v1 [q-bio.QM] 27 Jun 2016

I. INTRODUCTION

The accurate modelling and simulation of biological processes such as gene expression or signalling has gained a lot of interest over the last years and a large body of literature discussing various types of models, their inference and means of simulation has emerged. The main purpose of these models is to reflect and faithfully reproduce observed biological dynamics, while giving an insight into the underlying bio-molecular mechanisms.

One important aspect in the design of these models is the determination of the model parameters. Often there exists a mechanistic model of the cellular processes but parameters (e.g. reaction rates or initial molecule concentrations) are not known. Since the same network topology may result in different behaviour depending on the chosen parameters ([15]), this presents a major problem for model inference and explains the need for parameter estimation techniques.

The models used in Systems Biology can be coarsely classified into two groups: deterministic and stochastic models. Deterministic models usually rely on ordinary differential equations which, given the parameters and initial conditions, can describe the development of the biological system in a deterministic manner. However, many cellular processes like gene expression are subject to stochastic fluctuations. Traditionally such models have been treated deterministically, relying on the large copy numbers of the involved molecules to take care of any stochastic effects (see [17] for a review). It has been indicated though that the randomness of these systems may play an important role ([7], [19], [23]), stirring an increasing interest in stochastic models over the last years ([18], [33], [22], [14], [5]). Such stochastic models are usually described in the framework of stochastic chemical reaction networks that can be simulated using the Stochastic Simulation Algorithm (SSA) ([9]).

The inference of the model parameters is done from mainly two kinds of biological data. First, distribution data, which is data that reflects a whole population of cells at a particular time point and shows the distribution of the measured output in the entire population. The other is trajectory data, which focuses on single cells and monitors the behaviour throughout time. In the past years the availability of trajectory data has drastically increased, providing detailed information of the development of single cells throughout time. Unfortunately, while there has been plenty of research on parameter estimation for deterministic systems ([32], [25], [31], [4]) and some promising work on stochastic systems ([18], [28]), there is only very little literature available on stochastic systems using trajectory data, and even the available methods are computationally very demanding (see for instance [12], [11], [2], [29]). Further, these methods often rely on approximating the model dynamics (for instance using the diffusion approximation ([10]) or linear noise approximation ([6])), but these approximations may not always be justifiable (in the case of low copy numbers of the reactants for example) and might obscure crucial system behaviour.

One particular problem that is common to most inference methods is the usually high dimensional parameter space. Most of the sampling based inference techniques require the exploration of the full parameter space in order to find areas of "good" parameters, which is no easy task as the dimension of the parameter space increases.

We are proposing an inference method that relies on two main ideas. We utilize Dirichlet Process Gaussian Mixture Models for the estimation of the distribution of parameters (see for instance [13]). This technique is a sophisticated machine learning tool that shows very good performance even for high-dimensional spaces and is further very robust to outliers. We then employ a form of nested sampling, that iteratively samples parameter vectors from a sequence of distributions, that can be used to approximate the full posterior. The parameter vectors are sampled in an independent manner from the estimated densities making our method fully parallelizable.

In the following we introduce the model class we will be dealing with and give a brief overview over the task of

parameter inference. Then we illustrate the mechanisms of our method and demonstrate its performance on high-dimensional deterministic as well as stochastic problems.

A. Chemical Reaction Networks

We are considering a U-dimensional Markov Process $X(t)$ depending on a V-dimensional parameter vector θ. We denote with $X_u(\tau)$ the u^{th} entry of the state vector at time τ and with $X(\tau) = \{X_u(\tau)\}_{u=1,\dots,U}$ the state vector at time τ.

In the context of stochastic chemical reaction networks this Markov process describes the development of U species $\boldsymbol{X}_1, \boldsymbol{X}_2, \dots, \boldsymbol{X}_U$ through R reactions $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_R$ written as

$$\mathcal{R}_r = \sum_u^U p_{ru} \boldsymbol{X}_u \to \sum_u^U q_{ru} \boldsymbol{X}_u,$$

where p_{ru} and q_{ru} are the numbers of consumed and produced species for reaction r. The states $X_u(\tau)$ correspond to the number of molecules of species $\boldsymbol{X}_u$ at time τ and each reaction $\mathcal{R}_r$ has an associated propensity. The propensities as well as other properties of the model depend on a V-dimensional parameter vector θ.

The other modelling approach we consider is the deterministic, where the state at time τ follows a deterministic dynamic, described by an ordinary differential equation (ODE)

$$\frac{\mathrm{d}}{\mathrm{d}\tau} X(\tau) = f(X(\tau), \tau, \theta),$$

with a function f determined by the chemical reaction network. Such a formulation is particularly useful, when instead of single cell trajectories, one uses the average of many such trajectories, since the mean behaviour of a stochastic chemical reaction network can be described with an ODE.

B. General Task

Regardless of the chosen modelling, the process $X(t)$ is assumed to be not directly observable but only through a P-dimensional observation vector

$$Y_\tau \sim p(\cdot | X_\tau, \theta).$$

The observations Y_τ at time τ depend on the state X_τ and on the V-dimensional parameter vector θ.

It is assumed that the variable Y is not observed at all times but only on T timepoints $\tau_1, \dots, \tau_T$ and only for M different trajectories. The observed state for trajectory m and time point τ_t will be denoted with y_t^m. The m^{th} trajectory until the t^{th} time point will be denoted with $\boldsymbol{y}_t^m = \{y_1^m, \dots, y_t^m\}$ and with $\boldsymbol{y} = \{\boldsymbol{y}_T^m\}_{m=1,\dots,M}$ we denote all trajectories at all time points. If the observation is obtained by simulation using a certain parameter θ we will write $\boldsymbol{y}_\theta$.

In the Bayesian approach the parameter vector θ is treated as a random variable with associated prior $p_0(\theta)$. The goal is not to find just one set of parameters, but much rather to compute the posterior distribution $p(\theta|\boldsymbol{y})$ of θ

$$p(\theta|\boldsymbol{y}) \propto p(\boldsymbol{y}|\theta) p_0(\theta),$$

where $p(\boldsymbol{y}|\theta)$ is the likelihood of θ for the particular observation $\boldsymbol{y}$. This has several advantages over a single point estimate as it gives insight over the areas of "good" parameters as well as about their relevance for the simulation outcome (a wide posterior indicates non identifiability for example). The challenge often lies in the fact that for stochastic systems the likelihood $p(\boldsymbol{y}|\theta)$ is not easily available and needs to be approximated.

In this paper we follow the Bayesian approach and try to recover the posterior of the particle $p(\theta|\boldsymbol{y})$.

II. RESULTS

In the following we first outline the proposed algorithm to give a general idea of our approach, before we describe it in full detail.

Algorithm Outline

ALGORITHM 1. Algorithm Outline. An overview over the involved steps of our proposed algorithm. The plue part illustrates the part that is parallelized.

1: Given observations $\boldsymbol{y}$, a prior $p_0(\theta)$ for θ, and a sequence of thresholds ϵ_k
2: sample particles from p_0
3: **for** k=1, ... **do**
4: Approximate super level set of particles θ with $p(\boldsymbol{y}|\theta) \geq \epsilon_k$ using previous particles, DP-GMM and rejection sampling.
5: Independently sample new particles from approximated super level set.
6: **end for**

In this work we present a novel algorithm that reliably infers model parameters for large stochastic or deterministic models using trajectory data. We found a way to efficiently sample parameter vectors (particles) from the super level set of the likelihood (sets of particles with a likelihood equal to or higher than some threshold) corresponding to an increasing sequence of thresholds ϵ_k. From these samples from different super level sets we can then recover the full posterior (see Supplementary Material II). A brief outline of the algorithm can be seen in Algorithm 1 and a visualization fo the algorithm is shown in Figure 1.

We approximate these high dimensional super level sets through Dirichlet Process Gaussiam Mixture Models

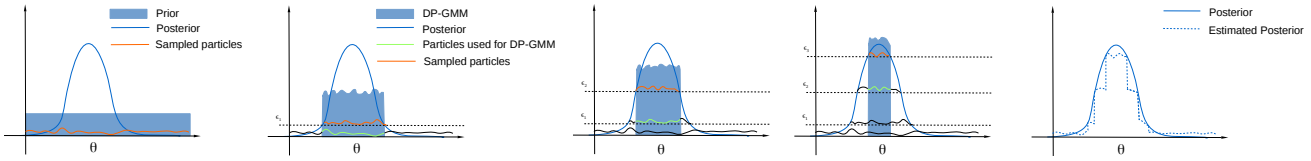

FIG. 1. Illustration of the algorithm. First particles (orange) are sampled from the prior (or previous DP-GMM estimation). Then these particles (those corresponding to a super level set of the likelihood, green) are being used to approximate the distribution of the super level set. Then the DP-GMM approximation of the super level set is used to sample a new set of particles corresponding to the super level set (orange). The sampling of the particles from the DP-GMM can be parallelized.

(DP-GMM) and rejection sampling (see Line 4 in Algorithm 1, see section IV). DP-GMM prove to be very efficient even with comparably few samples and are robust to outliers, which allows us to obtain a good approximation of the super level set, even when the likelihood of the underlying particles is not exactly known. The samples drawn for each super level set are independent (Line 5 in Algorithm 1), which allows us to parallelize the sampling of the particles. The sequential increase of the likelihood threshold allows us to efficiently explore the parameter space.

This exploration of the parameter space through super level sets of the likelihood is very similar to what is used in Nested Sampling approaches (as presented for instance in [1] or [24]), which we will briefly outline. Nested Sampling relies on iteratively sampling particles and computing their corresponding likelihoods. At first the samples are drawn from the prior. From the initially sampled particles the particle with the lowest likelihood is rejected and a new particle is sampled and accepted only if its likelihood is higher than the current lowest likelihood. This approach results in a sequence of samples from different super level sets, from which the posterior can be recovered. The usual formulation of Nested Sampling replaces in each iteration the particle with the lowest likelihood with a new particle, which makes it hard to parallelize this approach. Further, the sampling of the new particle is challenging and usually involves some volume estimation or random walk and the knowledge of the exact likelihood (see for instance [16],[20] or [8])

Since for stochastic models, the exact likelihood is usually not available, we rely on a particle filter approximation of the likelihood (see for details Supplementary Materials I). Such likelihood approximations are frequently used in Monte Carlo Markov Chain (MCMC) methods (see e.g. [11] or [5]). These methods construct samples from a Markov chain, whose stationary distribution is the posterior distribution of the parameters that we are interested in. MCMC methods are exact in the sense that they generate samples from the exact posterior (up to the inaccuracy of the likelihood approximation), but suffer from several known weaknesses, such as a possibly long burn in time, possibly high autocorrelation and thus poor exploration of the parameter space, no straight forward parallelization and required extensive tuning by the user.

Our proposed method is very well parallelizable and explores even high dimensional parameter spaces efficiently. In the following we elaborate on the details of the algorithm.

Algorithm Details

From now on we use the convention that for any variable p, $\hat{p}$ denotes its approximation. If p is a distribution this approximation is usually a Dirichlet Process Gaussian Mixture approximation (see section IV). This approximation corresponds to an approximation of the full distribution, while when l is the likelihood, the approximation $\hat{l}$ is usually obtained using a particle filter approach (see section Supplementary material I) and corresponds to an approximation of the likelihood evaluated at a single point.

ALGORITHM 2. Likelihood Nested Sampling Algorithm

1: Given observations $\boldsymbol{y}$, a prior $p_0(\theta)$ for θ, and some $0 < \alpha < 1$.
2: Set $\epsilon_1 = 0$, sample N particles from p_0 to obtain the set $\mathcal{S}_0$
3: **for** k=1, ... **do**
4: Estimate density $\hat{q}_{k-1}$ from the set $\mathcal{S}_{k-1}$
5: Set $\mathcal{S}'_k = \mathcal{S}_{k-1}$
6: Remove all particles θ from $\mathcal{S}'_k$ with $\hat{l}(\theta) < \epsilon_k$
7: Estimate density $\hat{q}'_k$ from the set $\mathcal{S}'_k$
8: **while** number of particles in $\mathcal{S}'_k$ is less than N **do**
9: sample $\theta \sim q'_k q_{k-1}^{-1} \approx L_k$.
10: **if** $\hat{l}(\theta) \geq \epsilon_k$ **then**
11: accept (θ_k^n) to $\mathcal{S}'_k$
12: **end if**
13: **end while**
14: Set $\mathcal{S}_k = \mathcal{S}'_k$
15: **end for**

The goal of the algorithm is to iteratively sample particles from the level set distributions

$$L_k(\theta) \propto \begin{cases} 1 & \text{if } l(\theta) > \epsilon_k \\ 0 & \text{otherwise} \end{cases}$$

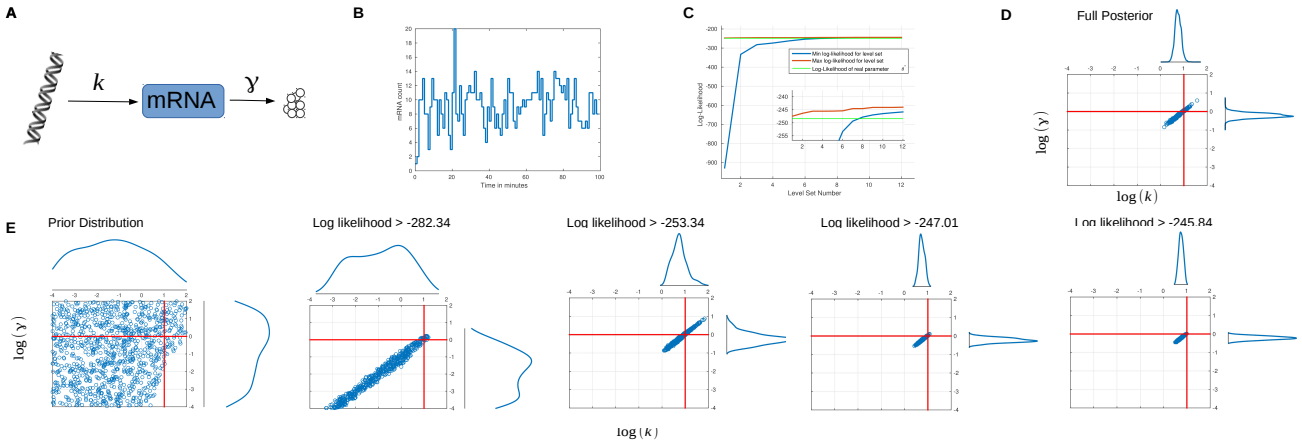

FIG. 2. **A**: Schematic of the birth-death model. mRNA gets transcribed from DNA and degrades with fixed rate constants. **B**: A single stochastically simulated trajectory of the birth-death system with birth rate $k = 10$ and degradation rate $\gamma = 1$. **C:** The development of the smallest (blue) and largest (red) log-likelihood of the level-set for each iteration. The green line indicates the log-likelihood of the true parameter θ^* that was used for the simulation.. **D**: The posterior of the log of the transcription and degradation rates as obtained with the LNS algorithm using only one simulated trajectory. The red lines indicate the true log-parameter values. **E**: The level sets of the log-likelihood of the log-translation and log-degradation rates as obtained with the LNS algorithm using only one observed trajectory.

for different ϵ_k. Since it is in general not possible to sample from the level sets directly, we create a sequence of sets $\mathcal{S}'_k$ which we use, combined with DP-GMM density approximation and rejection sampling, to sample from L_k. Once we have these level sets, we can recover the full posterior as described in the Supplementary material II. We first sample a set $\mathcal{S}_0$ of N particles

$$\mathcal{S}_0 = \{\theta_i^0\}_{i=1,\dots,N}$$

from the prior $p_0(\theta)$ and approximate for each θ_i the likelihood $l(\theta_i)$. The particles in set $\mathcal{S}_0$ are distributed according to the prior which we will denote with

$$\mathcal{S}_0 \sim p_0.$$

Our goal is to use this set to sample from the level set L_1. To this end we discard all particles with a likelihood lower than the threshold ϵ_1 (which we choose to be the α_1-quantile of the likelihoods in the set $\mathcal{S}_0$ for some α_1) to obtain a set of $N'_1 = \lfloor (1-\alpha_1)N \rfloor$ particles

$$\mathcal{S}'_1 = \{\theta_i^1 | \hat{l}(\theta_i^1) > \epsilon_1\}_{i=1,\dots,N'_1}.$$

The particles in this set are distributed according to $q'_1 := p_0(\theta)L_1(\theta)$

$$\mathcal{S}'_1 \sim \underbrace{p_0(\theta)L_1(\theta)}_{q'_1}.$$

This set can be used to approximate the distribution

$$q'_1 = p_0(\theta)L_1(\theta) \propto \begin{cases} p_0(\theta) & \text{if } \hat{l}(\theta) > \epsilon_1 \\ 0 & \text{otherwise} \end{cases}$$

by using DP-GMM (see IV). Now we can use rejection sampling to sample $\alpha_1 N$ particles from $L_1(\theta)$, by sampling θ^* from q'_1 and accepting it with a probability proportional to $p_0(\theta^*)^{-1}$. Together with the original N' particles we obtain the set

$$\mathcal{S}_1 \sim \underbrace{\frac{N'_1}{N}q'_1 + \frac{N - N'_1}{N}L_1}_{q_1}$$

of N particles with a likelihood larger than ϵ_1. Observe that this way we only have to sample $N - N'_1 = \lceil \alpha_1 N \rceil$ particles in each iteration which decreases the runtime drastically. We repeat this procedure iteratively. Assuming we have the set

$$\mathcal{S}_k \sim q_k,$$

we sample $\lceil \alpha_{k+1} N \rceil$ particles from the level set L_{k+1} by discarding all particles from set $\mathcal{S}_k$ with a likelihood lower than ϵ_{k+1} to obtain

$$\mathcal{S}'_{k+1} \sim \underbrace{\hat{q}_k L_{k+1}}_{q'_{k+1}},$$

use DP-GMM to obtain an approximation $\hat{q}'_{k+1}$ of q'_{k+1} and generate samples from the level set L_{k+1} by sampling θ^* from $\hat{q}'_{k+1}$ and accepting it with probability proportional to $\hat{q}_k(\theta^*)^{-1}$ ($\hat{q}_k$ is the DP-GMM approximation of q_k). The use of rejection sampling is crucial, since $\hat{q}_k$ as well as $\hat{q}'_{k+1}$ are DP-GMM approximations of sets of particles with only approximated likelihoods. Thus the rejection sampling accounts not only for a possible asymmetric prior, but also for any estimation error from the

density or likelihood approximation (see Supplementary Material III).

Consequently we also have

$$\mathcal{S}_{k+1} \sim \frac{N'_{k+1}}{N}\hat{q}'_{k+1} + \frac{N - N'_{k+1}}{N}L_{k+1}(\theta) =: q_{k+1}.$$

In the following we demonstrate its performance on three chosen examples.

The stochastic birth-death Model

We first demonstrate the basic mechanisms of the proposed method on the very simple birth-death model, illustrated in Figure 2 A. This example serves well as a first illustration since it only has two parameters. In this model a single species (mRNA) is constitutively generated at a birth rate k and degrades at a degradation rate γ. We simulate a single trajectory (Figure 2 B) using the Stochastic Simulation Algorithm ([9]) taking $k = 10$ and $\gamma = 1$ and an initial mRNA molecule number of 0. Using this simulated trajectory we infer the used parameters k and γ. We assume the initial mRNA number to be known. We assume a normal measurement noise with zero mean and a standard deviation of 2. At first our algorithm generates a set of parameter vectors (particles) that are sampled from a log-uniform distribution on the interval $[0.00001\ 100] \times [0.00001\ 100]$ and approximates the likelihood of each of those particles using a particle filtering approach (see Supplementary material I). If the likelihood could not be approximated (as for the particles in the lower right corner of the first graph of Figure 2 E), the particles are not accepted. This can happen if the sampled particles result in simulations so far from the observed data, that the likelihood is approximated with 0.

It then discards all particles with a likelihood smaller than a certain threshold which results in a discrete approximation of the super level set of the likelihood corresponding to that very threshold. By sampling particles from the current super level set and discarding all particles with a likelihood smaller than the next threshold, our algorithm generates a sequence of super level set distributions for the parameter (see Figure 2 E). If we are interested in only one parameter vector that results in simulations close to our observation, we can pick one of the particles from the highest level set, otherwise we can also recover the full posterior as illustrated in Figure 2 D. The resulting posterior is centred at the true parameter values, but its width indicates, that the parameters could not be identified uniquely, which is expected, since a wide range of parameters can result in the same trajectory. Figure 2 C shows the development of the log-likelihood threshold and the highest achieved log-likelihood for each level set. As can be seen, the final thresholds are even higher than the log-likelihood of the true parameter. This is also not surprising, since the used trajectory is only one realization of the stochastic systems and certain other parameters may be just as likely to have resulted in the same trajectory. We would like to point out that the super level sets illustrated in Figure 2 E seem to have an ellipsoid shape, as we would expect in this case.

A large stochastic model: The Lac-Gfp model

We demonstrate how our algorithm deals with a realistic sized stochastic model, by inferring the posterior for the parameters of the LacGfp model illustrated in Figure 3 A. This model has been already used in [18] as a benchmark, although with distribution-data (for details see Supplementary material and [18]). Here we use the model to simulate a number of trajectories and illustrate how our approach infers the posteriors of the used parameters. This model is particularly challenging in two aspects. First, the number of parameters is 18, making it a fairly large model to infer. Secondly, the model exhibits switch-like behaviour which makes it very hard to approximate the likelihood of such a switching trajectory (see the Supplementary material for a further discussion). In Figure 3 we show 10 of the simulated trajectories. As can be seen, most trajectories exhibit noise fluctuations first and at some point "switch on". The switching occurs randomly and the observed fluorescence increases rapidly in orders of magnitude. This raises considerable problems for the likelihood estimation. The measured species in this example is fluorescent GFP where it is assumed that each GFP-molecule emits fluorescence according to a normal distribution. We used 5 trajectories to infer the posterior, whose marginals are shown in Figure 3 C. As can be seen, the posterior seems to capture the real parameter fairly well. It is also seen that some parameters seem to have a wider marginal posterior than others. This indicates that the model is more sensitive to those parameters with narrower marginal posterior, since the likelihood seems to be more sensitive in respect of those parameters. The development of the log-likelihood threshold illustrated in Figure 3 D shows similar behaviour as for the birth-death model. The threshold eventually surpasses the actual log-likelihood of the real parameter θ^*. The first time that a particle is found with a higher log-likelihood than θ^* is after 30 iterations.

A realistic example: The deterministic Raf-Mek-Erk signalling cascade

Our method can similarly be applied to deterministic models. In a deterministic setting, the hidden system states X_τ do not develop according to a Markov Chain but according to a set of ordinary differential equations that depend on the parameter vector θ. Again, it is assumed that the states are not directly observable but only through noisy measurements $Y_\tau \sim p(\cdot|X_\tau, \theta)$. In the deterministic case the computation of the likelihood simplifies, since we do not need a particle filter but instead can solve the ordinary differential equations of the model with the particular parameters and compute the likelihood from the resulting trajectory. The use of DP-

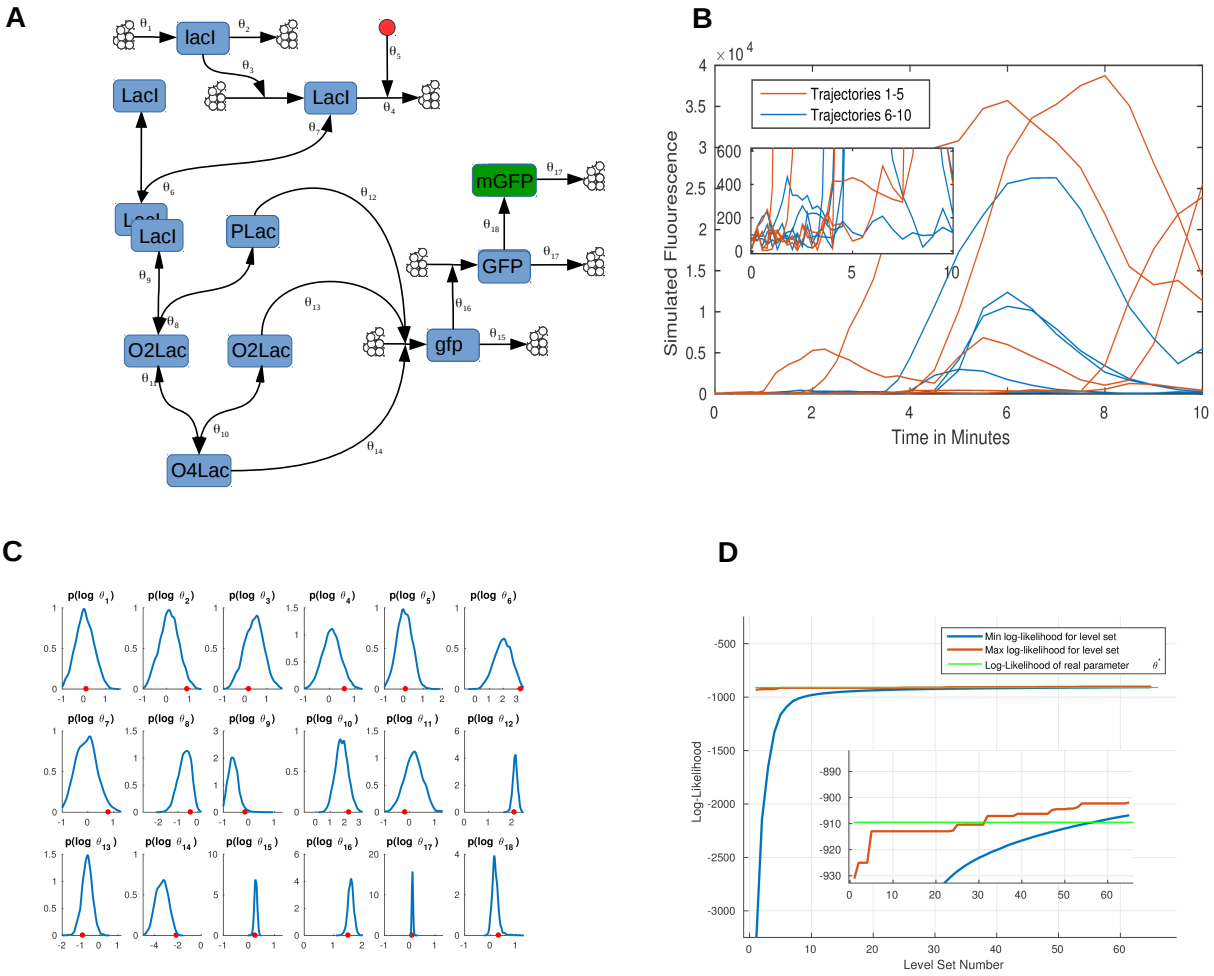

FIG. 3. **A**: The Lac-Gfp model, containing 9 species and 18 reactions. **B**: 10 Simulated model trajectories that were used for parameter inference. The system exhibits stochastic switch behaviour. **C**: Inferred posterior of the logarithm of the parameters. The red dots indicate the true parameter value θ^* that was used for the simulation of the data. As prior, a log-normal distribution in the indicated bounds was used. **D**: The development of the smallest (blue) and largest (red) log-likelihood of the level-set for each iteration. The green line indicates the log-likelihood of the true parameter θ^* that was used for the simulation.

GMM density estimation allows us to deal with high dimensional parameter spaces and the sequential nature of our algorithm makes a parallelization straight forward. We demonstrate our method on the example of the Ras-Raf-Mek-Erk pathway (Figure 4). We use the model introduced in [27], which consists of an activation cascade formed by the Egf receptor, Ras, Raf, Mek and Erk (see Figure 4 C). The model also features a slow negative feedback through Dusp and a fast receptor dependent negative feedback. The read out is the amount of phosphorylated Erk as measured by a FRET sensor. For details on the FRET sensor see [27]. As training data for our parameter inference we used the observed mean behaviour of the phosphorylated Erk upon different stimulations, see Figure 4 A. In order to check our inferred posterior, we picked the particle with the highest likelihood from the posterior and simulated the model for different stimuli. Figure 4 A shows the model behaviour for stimuli that have been used to identify the system, while Figure 4 B shows the model simulations for stimuli that have not been used for the inference. We see that using the inferred particle the model is able to reproduce the training data as well as predict the systems behaviour to new stimuli. We can use the inferred posterior to draw biological conclusions that, as a future step, could be verified experimentally. For instance, one of the model assumption is that dusp gets transcribed constitutively at a certain rate $\text{dusp}_{\text{basal}}$ and induced by Erk with the rate dusp_{ind}. Along all other parameters, both rates were inferred from the provided data as well as the dusp and DUSP initial concentrations. Their marginal posterior distribution is shown in Figure 5. From a quick

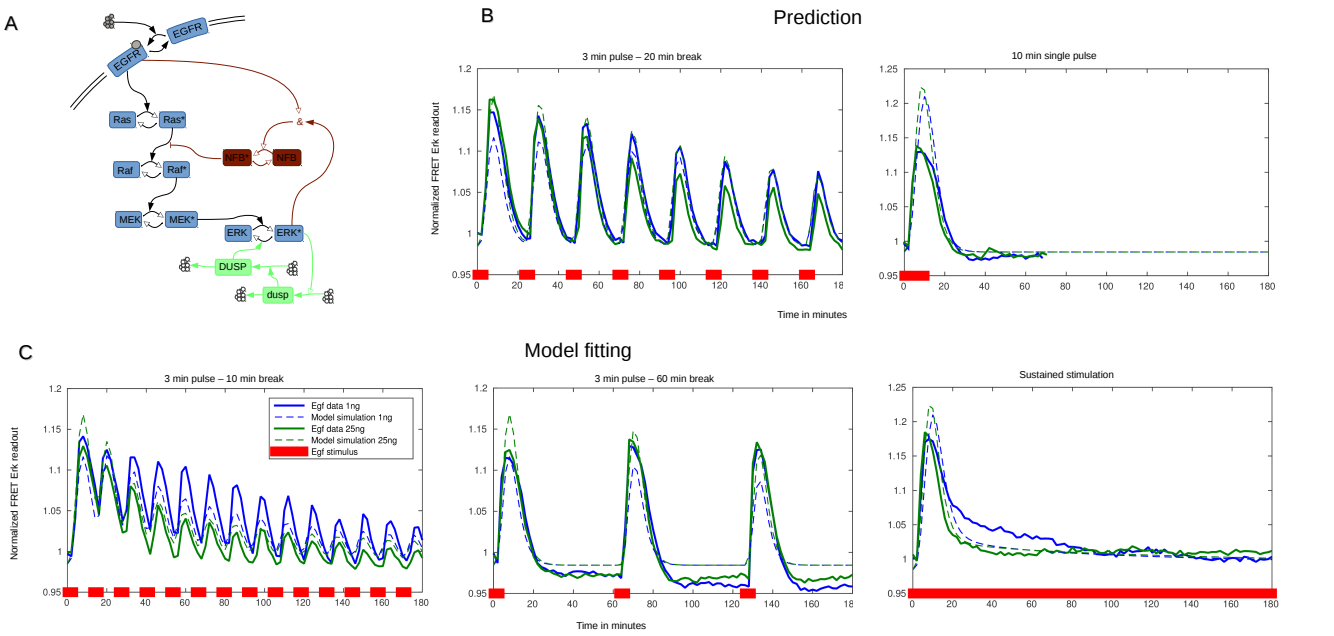

FIG. 4. **A**: Schematic representation of the Ras-Raf-Mek-Erk pathway model from [27]. Upon stimulation with Egf, the Egf-Receptor activates Ras, which in turn activates Raf, Mek and Erk. Dusp transcription is stimulated by phosphorilated Erk, while at the same time Dusp dephosphorilates Erk, making it a slow negative feedback from Erk upon itself. The NFB species is an unspecified, Erk and Egf dependent negative feedback species. The model has a total of 36 unknown parameters and 11 species. **B**: Prediction of the model. The particle with the highest likelihood was used to simulate an input stimulus, previously assumed unknown for the model inference. **C**: Fitting of the Egf-Ras-Raf-Mek-Erk model. The inference was performed using these 6 trajectories. The particle with the highest likelihood was then used for the simulation.

look at the marginal posteriors we see that in order for the model to produce the observed behaviour, the constitutive dusp transcription needs to be four orders of magnitude smaller than the Erk induced transcription. We further see that the initial concentration of dusp and DUSP needs to be very low ($\sim 10^{-3}$) compared with the initial concentration of Erk (0.6) for example. We mention these particular observations to illustrate that given the posterior, we can deduce biological statements that can be experimentally checked in order to investigate the validity of the model.

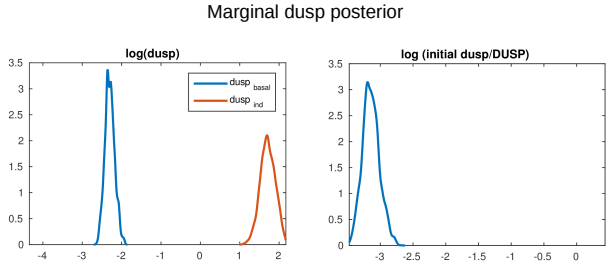

FIG. 5. The marginal posterior distribution for the Egf-Ras-Raf-Mek-Erk model of the constitutive transcription rate $\text{dustp}_{\text{basal}}$, the Erk induced dusp transcription dusp_{ind} and the initial concentration of dusp and DUSP (which is assumed to be equal).

III. CONCLUSION

We have introduced an inference method that reliably infers the posterior for model parameters of large stochastic and deterministic systems. We found a way to sample directly from the super level set of the likelihood without needing the exact likelihood and having to resort to volume approximations or random walks. This also allows us to parallelize the sampling of the super level sets. We have demonstrated the performance of our method applied to large stochastic models as well as realistic sized deterministic models and have shown how the obtained posterior of particles can be used not only to simulate and predict the model behaviour, but also to create biological statements that can be checked in order to verify model consistency.

IV. REMARKS

In the following we briefly outline some particular relevant details of the implementation of the algorithm.

Use quantiles for the computation of ϵ_{k+1}

To determine the sequence $\{\epsilon_k\}_k$, we use the α-quantile of q_k for some $0 < \alpha < 1$, where we choose the same α for all k,

$$\alpha_k = \alpha, \quad \forall k.$$

Thus from the set $\{\theta_i^k\}_{i=1,\dots,N}$ only $\lceil \alpha N \rceil$ particles will have a likelihood that is lower than ϵ_{k+1} and only these particles will need to be resampled.

DP-GMM for density estimation

In each iteration the density q'_k has to be approximated by $\hat{q}'_k$ from a set of N particles $\{\theta_i^k\}_{i=1,\dots,N}$. This is a task that needs to be performed for most Monte Carlo inference algorithms and is usually done with kernel density estimation (KDE) which places a perturbation kernel on each particle and takes $\hat{q}'_k$ to be the sum of all these kernels. This approach has several drawbacks as it is not clear what kernel to use or how to choose its width. Further, KDE is sensitive to outliers and shows a poor performance with sparse data. This poses a serious problem for most sample based inference techniques, since in most applications the parameter space is large and only comparably few particles are available.

To overcome these problems we propose to use a density estimation technique that approximates the target distribution with a mixture of Gaussians, where the number, shape and weight of each Gaussian is inferred from the particle population. The estimation technique - Dirichlet Process Gaussian Mixture Model (DP-GMM) ([13]) - uses a hierarchical prior on the mixture model and assumes that the mixture components are distributed according to a Dirichlet Process. The inference of the distribution is an iterative process that uses Gibbs sampling to infer the parameters and hyper parameters of the mixture model. This procedure is usually computationally demanding. However, since only $\lceil \alpha N \rceil$ particles are resampled in each population, the previous parameters and hyper parameters provide an already good prior estimate for the new estimation which allows to find a new estimation with only few iterations. DP-GMM estimations perform comparably well with sparse and high dimensional data, because they use the joint information of all samples to infer the high dimensional density. The use of DP-GMM is what enables us to apply our algorithm even to high dimensional parameter spaces. It also allows us to follow a nested sampling approach without needing the exact likelihood. Even when we assume that we do have the exact likelihood, which would reduce our algorithm to the already established nested sampling method, the use of the DP-GMM estimation allows us to efficiently sample from a high dimensional level set, without having to approximate the level set directly.

Further details are given in the Supplementary material and can also be found in [13], [26] and [30] (in particular [13] shows a comparison between DP-GMM and KDE).

[1] Stuart Aitken and Ozgur E Akman. Nested sampling for parameter inference in systems biology: application to an exemplar circadian model. *BMC systems biology*, 7(1):72, 2013.

[2] Aleksandr Andreychenko, Linar Mikeev, David Spieler, and Verena Wolf. Parameter identification for markov models of biochemical reactions. In *Computer Aided Verification*, pages 83–98. Springer, 2011.

[3] Marc R Birtwistle, Alexander von Kriegsheim, Katarzyna Kida, Juliane P Schwarz, Kurt I Anderson, and Walter Kolch. Linear approaches to intramolecular förster resonance energy transfer probe measurements for quantitative modeling. 2011.

[4] DM Bortz and PW Nelson. Model selection and mixed-effects modeling of hiv infection dynamics. *Bulletin of mathematical biology*, 68(8):2005–2025, 2006.

[5] Richard J Boys, Darren J Wilkinson, and Thomas BL Kirkwood. Bayesian inference for a discretely observed stochastic kinetic model. *Statistics and Computing*, 18(2):125–135, 2008.

[6] Johan Elf and Måns Ehrenberg. Fast evaluation of fluctuations in biochemical networks with the linear noise approximation. *Genome research*, 13(11):2475–2484, 2003.

[7] Michael B Elowitz, Arnold J Levine, Eric D Siggia, and Peter S Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.

[8] F Feroz, MP Hobson, E Cameron, and AN Pettitt. Importance nested sampling and the multinest algorithm. *arXiv preprint arXiv:1306.2144*, 2013.

[9] Daniel T Gillespie. Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry*, 81(25):2340–2361, 1977.

[10] Daniel T Gillespie. The chemical langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, 2000.

[11] Andrew Golightly and Darren J Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle markov chain monte carlo. *Interface Focus*, 1(6):807–820, 2011.

[12] Andrew Golightly and Darren J Wilkinson. Bayesian inference for markov jump processes with informative observations. *arXiv preprint arXiv:1409.4362*, 2014.

[13] Dilan Görür and Carl Edward Rasmussen. Dirichlet process gaussian mixture models: Choice of the base distribution. *Journal of Computer Science and Technology*, 25(4):653–664, 2010.

[14] Andreas Hilfinger and Johan Paulsson. Separating intrinsic from extrinsic fluctuations in dynamic biological systems. *Proceedings of the National Academy of Sciences*, 108(29):12167–12172, 2011.

[15] Piers J Ingram, Michael PH Stumpf, and Jaroslav Stark. Network motifs: structure does not determine function. *BMC genomics*, 7(1):108, 2006.

[16] Rob Johnson, Paul Kirk, and Michael PH Stumpf. Sysbions: nested sampling for systems biology. *Bioinformat-*

ics, 31(4):604–605, 2015.

[17] Guy Karlebach and Ron Shamir. Modelling and analysis of gene regulatory networks. *Nature Reviews Molecular Cell Biology*, 9(10):770–780, 2008.

[18] Gabriele Lillacci and Mustafa Khammash. The signal within the noise: efficient inference of stochastic gene regulation models using fluorescence histograms and stochastic simulations. *Bioinformatics*, 29(18):2311–2319, 2013.

[19] Harley H McAdams and Adam Arkin. Stochastic mechanisms in gene expression. *Proceedings of the National Academy of Sciences*, 94(3):814–819, 1997.

[20] Pia Mukherjee, David Parkinson, and Andrew R Liddle. A nested sampling algorithm for cosmological model selection. *The Astrophysical Journal Letters*, 638(2):L51, 2006.

[21] Brian Munsky and Mustafa Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of chemical physics*, 124(4):044104, 2006.

[22] Gregor Neuert, Brian Munsky, Rui Zhen Tan, Leonid Teytelman, Mustafa Khammash, and Alexander van Oudenaarden. Systematic identification of signal-activated stochastic gene regulation. *Science*, 339(6119):584–587, 2013.

[23] Ertugrul M Ozbudak, Mukund Thattai, Iren Kurtser, Alan D Grossman, and Alexander van Oudenaarden. Regulation of noise in the expression of a single gene. *Nature genetics*, 31(1):69–73, 2002.

[24] Nick Pullen and Richard J Morris. Bayesian model comparison and parameter inference in systems biology using nested sampling. *PloS one*, 9(2):e88419, 2014.

[25] Hein Putter, SH Heisterkamp, JMA Lange, and F De Wolf. A bayesian approach to parameter estimation in hiv dynamical models. *Statistics in medicine*, 21(15):2199–2214, 2002.

[26] Carl Edward Rasmussen. The infinite gaussian mixture model. In *NIPS*, volume 12, pages 554–560, 1999.

[27] Hyunryul Ryu, Minhwan Chung, Maciej Dobrzyński, Dirk Fey, Yannick Blum, Sung Sik Lee, Matthias Peter, Boris N Kholodenko, Noo Li Jeon, and Olivier Pertz. Frequency modulation of erk activation dynamics rewires cell fate. *Molecular Systems Biology*, 11(11):838, 2015.

[28] Scott A Sisson, Yanan Fan, and Mark M Tanaka. Sequential monte carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765, 2007.

[29] Vassilios Stathopoulos and Mark A Girolami. Markov chain monte carlo inference for markov jump processes via the linear noise approximation. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1984):20110541, 2013.

[30] Yee Whye Teh, Michael I Jordan, Matthew J Beal, and David M Blei. Hierarchical dirichlet processes. *Journal of the american statistical association*, 101(476), 2006.

[31] Jens Timmer, TG Müller, I Swameye, O Sandra, and Ursula Klingmüller. Modeling the nonlinear dynamics of cellular signal transduction. *International Journal of Bifurcation and Chaos*, 14(06):2069–2079, 2004.

[32] Tina Toni, David Welch, Natalja Strelkowa, Andreas Ipsen, and Michael PH Stumpf. Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6(31):187–202, 2009.

[33] Christoph Zechner, Michael Unger, Serge Pelet, Matthias Peter, and Heinz Koeppl. Scalable inference of heterogeneous reaction kinetics from pooled single-cell recordings. *Nature methods*, 11(2):197–202, 2014.

Supplementary

V. COMPUTING THE LIKELIHOOD

A. Computation of the likelihood

The most critical part in the algorithm is the computation of the likelihood, which is why we will quickly discuss one of the main issues and our approach of solving it. We follow the particle filter approach from [11], which we will briefly outline.

The likelihood for the set of trajectories $\{\boldsymbol{y}^m\}_{m=1\ldots,M}$ is computed as the product of the likelihoods of each trajectory

$$p(\boldsymbol{y}|\theta) = \prod_{m=1}^{M} p(\boldsymbol{y}^m|\theta).$$

The likelihood for each trajectory $\boldsymbol{y}^m$ is computed as the product of the likelihoods at each timepoint

$$p(\boldsymbol{y}^m|\theta) = \prod_{t=1}^{T} p(y_t^m|\boldsymbol{y}_{t-1}^m\theta).$$

To be precise, the likelihood for each trajectory up until timepoint $t+1$, $\boldsymbol{y}_{t+1}^m$ gets computed recursively

$$\hat{p}(\boldsymbol{y}_{t+1}^m|\theta) \approx \hat{p}(\boldsymbol{y}_t^m|\theta) \int p(y_{t+1}^m|X_{t+1})p(X_{t+1}|\boldsymbol{y}_t^m,\theta)dX_{t+1}.$$

The integral in the recursive formula is approximated by forward simulation of H hidden trajectories which are afterwards resampled. In particular, let us assume we have an empirical distribution of N states from the initial state distribution $X_0^i \sim p_{X_0}$, $i = 1, \ldots, N$. To obtain an approximation of the distribution $p(X_1|\boldsymbol{y}_0)$, we weight each of the H states X_0^h with the likelihood $w_0^h = p(\boldsymbol{y}_0|X_0^h)$ and resample them using those weights to obtain an empirical approximation for the distribution $p(X_0|\boldsymbol{y}_0, \theta)$. By forward simulating from X_0^h, we thus get the states $X_1^h \sim p(X_1|\boldsymbol{y}_0, \theta)$. The likelihood for $\boldsymbol{y}_1$ can then be approximated by

$$p(\boldsymbol{y}_1|\theta) = p(\boldsymbol{y}_0|\theta)\frac{1}{H}\sum_h p(\boldsymbol{y}_1|X_1^h). \tag{1}$$

This procedure of weighting, resampling and forward simulating allows us to approximate for each t the distribution $p(X_{t+1}|\boldsymbol{y}_t^m, \theta)$. Once this is accomplished, the computation of $p(y_{t+1}^m|X_{t+1})$ is straight forward. This procedure results in Algorithm 3.

Given H states X_0^i sampled from some initial distribution p_{X_0} and a parameter vector θ.
Compute for each i the likelihood $l_0^i = p(\boldsymbol{y}_0|X_0^h)$.
Set $p(\boldsymbol{y}_0|\theta) = \frac{1}{H}\sum_{i=1}^H l_0^i$.
Compute weights $w_0^i = \frac{l_0^i}{\sum_{i=1}^H l_0^i}$.
for t=1, …, T **do**
  Sample H indices i' according to weights w_{t-1}^i.
  Simulate X_t^i starting from $X_{t-1}^{i'}$ and using θ.
  Compute $l_t^i = p(y_t|X_t^i, \theta)$.
  Compute $p(\boldsymbol{y}_t|\theta) = p(\boldsymbol{y}_{t-1}|\theta)\frac{1}{H}\sum_{i=1}^H l_t^i$.
  Compute weights $w_0^i = \frac{l_0^i}{\sum_{i=1}^H l_0^i}$.
end for

ALGORITHM 3. Likelihood Approximation using particle filter

VI. APPROXIMATING THE FULL POSTERIOR

In Bayesian analysis the goal is to infer the full posterior $p(\theta|\boldsymbol{y})$ and not just the level sets $L_{\epsilon_k}(\theta)$. One obvious way of recovering the full posterior is to divide the parameter space into bins and take the average of all computed likelihoods in that bin. The number of available particles will be in general large, since our algorithm relies on sampling a large number of particles. This approach may, however, result in a very noisy approximation of the posterior, since the likelihoods are only approximated and gets quickly unwieldly as dimensions increase. We propose a different approach of recovering the likelihood by exploiting the approximation of the level sets $L_{\epsilon_k}(\theta)$.

A. Recovering the posterior from level sets

At first we assume we have exact level set distributions for the thresholds ϵ_k

$$L_k(\theta) \propto \mathbb{1}_{p(\boldsymbol{y}|\theta)\geq\epsilon_k}.$$

Considering the normalization constant we have the equality

$$L_k(\theta) = \frac{\mathbb{1}_{p(\boldsymbol{y}|\theta)>\epsilon_k}}{\int\limits_\Omega \mathbb{1}_{p(\boldsymbol{y}|\theta)>\epsilon_k} d\theta}.$$

Writing $g(\epsilon_k) = \int\limits_\Omega \mathbb{1}_{p(\boldsymbol{y}|\theta)>\epsilon_k} d\theta$, we can obtain the likelihood from the level sets by

$$\int\limits_0^\infty L_k(\theta)g(\epsilon_k)d\epsilon_k = \int\limits_0^\infty \mathbb{1}_{p(\boldsymbol{y}|\theta)>\epsilon_k}(\epsilon_k)d\epsilon_k = \int\limits_0^{p(\boldsymbol{y}|\theta)} d\epsilon_k = p(\boldsymbol{y}|\theta).$$

We can multiply this likelihood by the prior $p(\theta)$ to obtain the posterior

$$p(\theta|\boldsymbol{y}) \propto p(\boldsymbol{y}|\theta)p(\theta).$$

TABLE I. Species and initial numbers for the RNA model

Species	Notation	Initial Distribution
Gene in on state	G_{on}	fixed to 0
Gene in off state	G_{off}	fixed to 1
mRNA	RNA	fixed to 0

B. Use approximated level sets

To obtain an approximation of the function $L_k(\theta)$, we first need an approximation $\hat{p}(\boldsymbol{y}|\theta)$ for the likelihood $p(\boldsymbol{y}|\theta)$ in order to obtain a set of particles $\{\theta_i|\hat{p}(\boldsymbol{y}|\theta_i) > \epsilon_k)\}_{i=1,\dots,N}$. This set of particles is then used to obtain an approximation $\hat{L}_k(\theta)$. In V we explained how we obtain $\hat{p}(\boldsymbol{y}|\theta)$. The posterior can then be obtained through the approximation of the integral

$$\hat{p}(\theta|\boldsymbol{y}) = \sum_{k=1}^{K} \hat{L}_k(\theta)p(\theta)\hat{g}(\epsilon_k)\delta_{\epsilon_k} \approx \int L_k(\theta)g(\epsilon_k)p(\theta)d\epsilon_k,$$

where

$$\delta_{\epsilon_k} = \begin{cases} \epsilon_k & \text{if } k = 1 \\ \epsilon_k - \epsilon_{k-1} & \text{otherwise} \end{cases}.$$

We approximate $g(\epsilon_k)$ with $\hat{g}$ which we assume to be some appropriate approximation. We want $\hat{g}$ to be proportional to $\int_\Omega \mathbb{1}_{p(\boldsymbol{y}|\theta)>\epsilon_k} d\theta$. Since the volume of the confidence ellipsoid of the data is proportional to the product of the square roots of the eigenvalues of the covariance matrix of the data, we choose

$$\hat{g} = \sqrt{\det(\mathrm{Cov}L_k)}$$

where $\mathrm{Cov}L_k$ is the Covariance matrix of the distribution L_k. We like to note that in the case where the level sets are expected to be multi modal, this approximation may not be the best choice and one has to thing of different ways of approximating g, however in our encountered examples $\sqrt{\det(\mathrm{Cov}L_k)}$ proved to be a sufficiently good approximation.

C. Sampling from the posterior

We described how, in the course of the algorithm run, we sample from approximations of each level set $\hat{L}_k$. Since we can sample from the level set $\hat{L}_k$, we can also sample from the weighted sum

$$\sum_{k=1}^{K} \hat{L}_k(\theta)\hat{g}(\epsilon_k)\delta_{\epsilon_k}$$

and weighting each sample θ with $p(\theta)$ to obtain weighted samples form the posterior.

D. Verification of approximated posterior

To illustrate the validity of the recovered posterior, we consider a very simple RNA transcription model consisting of the species in table VI D.

The reactions read as follows:

Reactions

1. $G_{\text{off}} \xrightarrow{\theta_1} G_{\text{on}}$. Gene switch on.
2. $G_{\text{on}} \xrightarrow{\theta_2} G_{\text{off}}$. Gene switch off.
3. $G_{\text{on}} \xrightarrow{\theta_3} G_{\text{on}} + RNA$. Transcription of mRNA.

TABLE II. Prior distributions of the RNA-system parameters and the real values $\hat{\boldsymbol{\theta}}$ used for the simulation of $\boldsymbol{y}$.

Parameter	Meaning	Prior interval on θ_i	$\hat{\boldsymbol{\theta}}$
θ_1	Rate of switching Gene on	[0.000001, 0.006]	0.001
θ_2	Rate of switching Gene off	assumed known	0.002
θ_3	mRNA transcription rate	assumed known	0.06
θ_4	mRNA degradation rate	assumed known	0.002

4. $RNA \xrightarrow{\theta_4} \emptyset$. mRNA degradation.

The parameters with corresponding priors and values used for the simulation are shown in Table II. We used this system to simulate one trajectory, observing the development of mRNA. The trajectory was observed on 25 equally spaced timepoints between 0 and 200 time units. We used the FSP [21] to compute the (almost) actual posterior distribution of θ_1. We used this (almost) exact likelihood to perform the parameter inference and show in Figure **??** the histogram of the recovered posterior, where the red dot indicates to position of the real parameter. As can be seen the posterior is nicely recovered from the level sets.

VII. REJECTION SAMPLING

In the process of the algorithm we need to sample particles from the level set L_k using the set

$$\mathcal{S}'_k \sim \underbrace{q_{k-1}L_k}_{q'_k}$$

to construct the set $\mathcal{S}_k$. We have no means to sample directly from L_k, but we have the Dirichlet Process Gaussian Mixture Model density estimations $\hat{q}'_k$ and $\hat{q}_{k-1}$ of the densities q'_k and q_{k-1} respectively. We generate samples from L_k by drawing a sample θ^* from $\hat{q}'_k$ and accept it with a probability $\frac{M_k}{q_{k-1}(\theta^*)}$ with an appropriately chosen M_k. This guarantees that the sampled particle θ^* is distributed according to

$$\theta^* \sim \frac{\hat{q}'_k M_k}{\hat{q}_{k-1}} \approx M_k L_k$$

(with a slight abuse of notation since the right hand side is not normalized). Since we only have approximations of q'_k and q_{k-1} and we cannot expect their ratio to be the actual level set, we also approximate the likelihood $l(\theta^*)$ and reject the particle if the approximated likelihood $\hat{l}(\theta^*)$ is smaller than the threshold ϵ_k. The sampling of the particles is a crucial step in the LNS algorithm, which is why we will discuss the two important parts of our rejection sampling approach

- Choice of M_k
- DP-GMM approximation

A. Choice of M_k

This rejection sampling approach is guaranteed to give us the right target distribution L_k whenever $M < q_{k-1}(\theta^*)$. Thus M determines the domain from which we want to sample. If the particle θ^* (that is sampled from $\hat{q}'_k$) is outside of this domain, we reject it automatically. For numerical reasons it is favourable to pick M_k not too small. Particularly for θ where $q'_k(\theta) = q_{k-1}(\theta) = 0$ but $\hat{q}'_k(\theta) \neq 0 \neq \hat{q}_{k-1}(\theta)$, choosing M_k too small may cause large numerical errors. In practice we did not observe a huge influence of the value for M_k on the algorithm performance as long as it was smaller than $\max_\theta \hat{q}_{k-1}$. For our algorithm runs we chose M_k to be the value of the 5% quantile of $\hat{p}_{k-1}$. The main function of the rejection sampling is guaranteeing that estimation errors of $\hat{q}'_k$ and $\hat{q}_{k-1}$ don't propagate through iterations. For an illustration on the RNA model see Figure **??**.

B. DP-GMM Approximation

The distributions q'_k and q_{k-1} are approximated using Dirichlet Process Gaussian Mixture Model approximation. This means that each of the distributions is approximated using Gaussian mixtures. This may at first glance not seem

like an appropriate way of approximating level sets due to their non-smooth nature. However, DP-GMM estimations work very well in high dimensional spaces and since we use rejection sampling to obtain samples from the level set by sampling from the DP-GMM estimation, the estimation error does not get propagated through iterations.

VIII. TERMINATION CRITERION

In the presentation of the algorithm we did not provide a termination criterion. An investigation of the development of the minimal and maximal log-likelihood threshold of inference runs for one trajectory and different values for α and different number of particles for the likelihood computation suggest two possible termination criteria.

- **The difference between minimal and maximal log-likelihood for a level set**
  As the algorithm proceeds the difference between the maximal and the minimal log-likelihood for each level set decreases (see for instance Figure 10 B or ...). In general we cannot expect this difference to become 0 due to the fact that we always deal with an approximation of the likelihood. One way of defining a termination criterion is to define a threshold and as soon as the difference of the maximal and minimal log-likelihood is smaller than the predefined threshold the algorithm will stop. The difference of the maximal and minimal likelihood seems to be a monotonically decreasing function. Assuming that there exists a maximal log-likelihood on the parameter space, the difference is guaranteed to approach 0 (but maybe not monotonically). This way the algorithm is guaranteed to terminate.

- **A low acceptance rate**
  It is an interesting observation, that the acceptance rate for each iteration drops as the algorithm progresses. This brings us to the question how we would expect the acceptance rate to behave. At iteration k the new particle θ^* is sampled from the approximative distribution

  $$\theta^* \sim \hat{L}_k(\cdot) \approx L_k(\cdot) = \mathbb{1}_{p(\boldsymbol{y}|\theta^*)>\epsilon_k}(\cdot)$$

  and is accepted if $\hat{p}(\boldsymbol{y}|\theta^*) > \epsilon_k$. Thus we can expect that most particles will be accepted and that the acceptance rate depends on the accuracy of the approximations $\hat{L}_k$ and $\hat{p}(\boldsymbol{y}|\theta^*)$. In practise the approximation of $\hat{L}_k$ is done by DP-GMM and rejection sampling, whose accuracy does not depend on the iteration number. This leaves the accuracy of the likelihood approximation as the main contributor for the change in acceptance rate. More precise, the acceptance rate depends on the quotient of the distance of the likelihood of θ^* to the threshold ϵ_k and the variance of the likelihood estimation $\hat{p}(\boldsymbol{y}|\theta^*)$

  $$\frac{p(\boldsymbol{y}|\theta^*) - \epsilon_k}{\text{var } \hat{p}(\boldsymbol{y}|\theta^*)}.$$

  In Figure 6 B, the acceptance rate for each log-likelihood threshold is plotted for different numbers of particles used for the likelihood approximation. As can be seen, the higher the number of particles (and thus the accuracy of the log-likelihood approximation) the higher is the acceptance rate.

  Thus the acceptance rate depends on the variance of the likelihood estimation $\hat{p}(\boldsymbol{y}|\theta^*)$, which can be assumed to be independent of particular iteration number k and on the average distance of the sampled particles θ to the current threshold ϵ_k. If the acceptance rate drops drastically, it can therefore be assumed that all particles are very close to the current threshold, which is an indicator for the local optimality. An illustration of the dropping acceptance rate is seen in Figure 6 where the acceptance rate is plotted against the threshold ϵ_k for different values of α (**A**) and for different number of particles for the log-likelihood approximation (**B**).

IX. EXAMPLES USED

A. Egf-Ras-Raf-Mek-Erk model

The Egf-Ras-Raf-Mek pathway is a signalling pathway, activated by Egf at the membrane and signalling through Ras, Raf and Mek to Erk (see Figure 7 B for an illustration). We use the model introduced in [27], with a slow negative feedback on Erk through Dusp and a fast, Egf-receptor dependent, negative feedback through a modelling species named "NFB". We use initial conditions different than in [27]. The involved species and corresponding initial concentrations can be taken from Table IX A. The fixed values were taken from the literature. The initial concentration of the active form of the proteins is assumed 0, since the cells were starved before the experiment. The

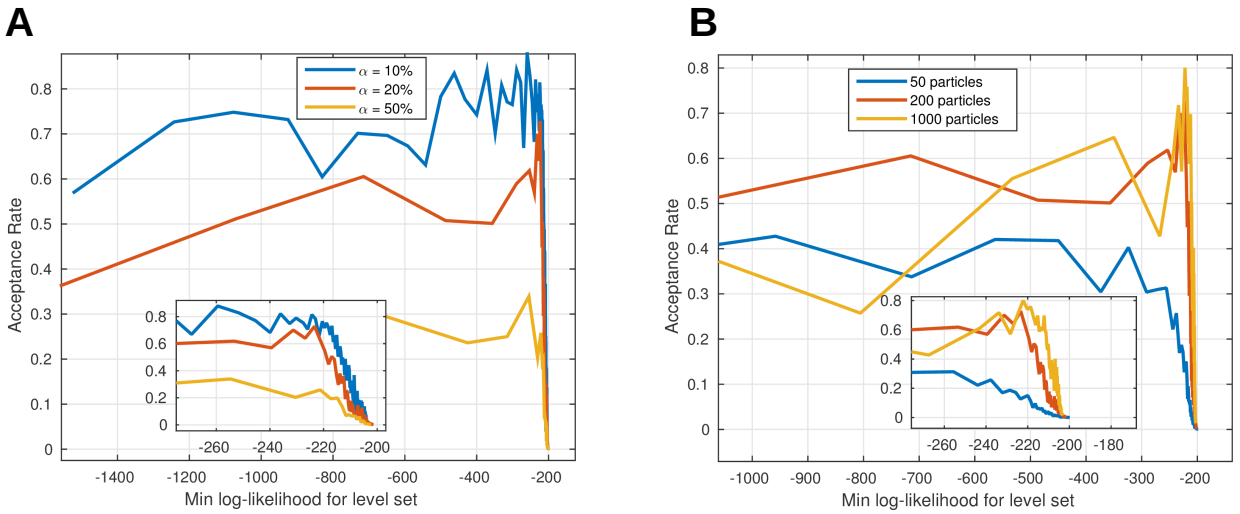

FIG. 6. **A:** Acceptance rate for each log-likelihood threshold for runs with different values for α. **B:** Acceptance rate for each log-likelihood threshold for runs with different numbers or particles for the likelihood approximation.

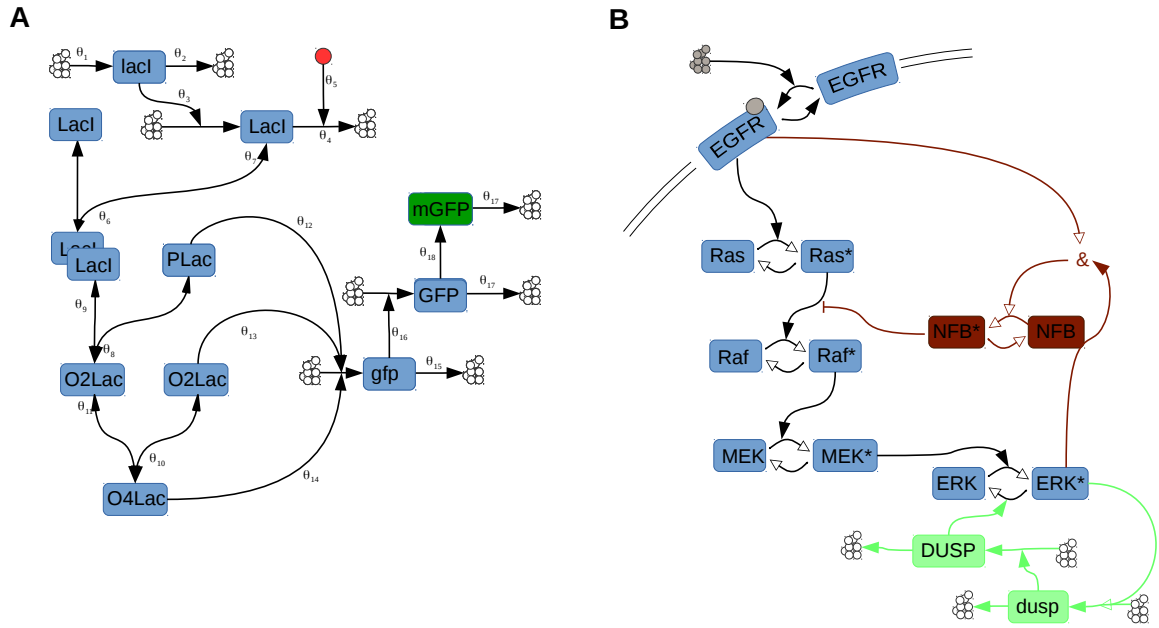

FIG. 7. **A**: Schematic representation of the LacGfp model. **B**: Schematic representation of the Egf-Ras-Raf-Mek-Erk model.

other values were taken from http://bionumbers.hms.harvard.edu/. For Dusp it was assumed that at the beginning of the experiment the concentration of mRNA and Protein was the same.

The reaction equations take the following form:

TABLE III. Species and intial concentration of the Ras-Raf-MEK-ERK model.

Species	Notation	Initial value prior interval / fixed value
Egf Receptor	EGFR	[exp(−8), exp(1)]
Active Egf Receptor	$EGFR^*$	fixed to 0
Ras concentration	Ras	fixed to 0.1
Active Ras concentration	Ras^*	fixed to 0
Raf concentration	Raf	fixed to 0.7 (b-Raf: 0.2, c-Raf: 0.5)
Active Raf concentration	Raf^*	fixed to 0
MEK concentration	MEK	fixed to 0.68
Active MEK concentration	MEK^*	fixed to 0
ERK concentration	ERK	fixed to 0.26
Active ERK concentration	ERK^*	fixed to 0
NFB concentration	NFB	[exp(−8), exp(1)]
Active NFB concentration	NFB^*	fixed to 0
Dusp mRNA/ Dusp Protein concentration	dusp/DUSP	[exp(−8), exp(1)]

1. *Reaction Equations*

$$\dot{\mathrm{EGFR}}(t) = -(k_1 \mathrm{egf}_{input}(t)\mathrm{EGFR}(t)) + (\gamma_1 \mathrm{EGFR}^*(t))$$

$$\dot{\mathrm{EGFR}}^*(t) = (k_1 \mathrm{egf}_{input}(t)\mathrm{EGFR}(t)) - (\gamma_1 \mathrm{EGFR}^*(t))$$

$$\dot{\mathrm{Ras}}(t) = -\left(k_2 \mathrm{EGFR}^*(t)\frac{\mathrm{Ras}(t)}{K_2 + \mathrm{Ras}(t)}\right) + \left(\gamma_2 \frac{\mathrm{Ras}^*(t)}{D_2 + \mathrm{Ras}^*(t)}\right)$$

$$\dot{\mathrm{Ras}}^*(t) = \left(k_2 \mathrm{EGFR}^*(t)\frac{\mathrm{Ras}(t)}{K_2 + \mathrm{Ras}(t)}\right) - \left(\gamma_2 \frac{\mathrm{Ras}^*(t)}{D_2 + \mathrm{Ras}^*(t)}\right)$$

$$\dot{\mathrm{Raf}}(t) = -\left(k_3 \mathrm{Ras}^*(t)\frac{\mathrm{Raf}(t)}{K_3 + \mathrm{Raf}(t)}\frac{\mathrm{K}_{NFB}^2}{\mathrm{K}_{NFB}^2 + \mathrm{NFB}^*(t)^2}\right) + \left(\gamma_3 \frac{\mathrm{Raf}^*(t)}{D_3 + \mathrm{Raf}^*}\right)$$

$$\dot{\mathrm{Raf}}^*(t) = \left(k_3 \mathrm{Ras}^*(t)\frac{\mathrm{Raf}(t)}{K_3 + \mathrm{Raf}(t)}\frac{\mathrm{K}_{NFB}^2}{\mathrm{K}_{NFB}^2 + \mathrm{NFB}^*(t)^2}\right) - \left(\gamma_3 \frac{\mathrm{Raf}^*(t)}{D_3 + \mathrm{Raf}^*}\right)$$

$$\dot{\mathrm{MEK}}(t) = -\left(k_4 \mathrm{Raf}^*(t)\frac{\mathrm{MEK}(t)}{K_4 + \mathrm{MEK}(t)}\right) + \left(\gamma_4 \frac{\mathrm{MEK}^*(t)}{D_4 + \mathrm{MEK}^*(t)}\right)$$

$$\dot{\mathrm{MEK}}^*(t) = \left(k_4 \mathrm{Raf}^*(t)\frac{\mathrm{MEK}(t)}{K_4 + \mathrm{MEK}(t)}\right) - \left(\gamma_4 \frac{\mathrm{MEK}^*(t)}{D_4 + \mathrm{MEK}^*(t)}\right)$$

$$\dot{\mathrm{ERK}}(t) = -\left(k_5 \mathrm{MEK}^*(t)\frac{\mathrm{ERK}(t)}{K_5 + \mathrm{ERK}(t)}\right) + \left(\gamma_5 \mathrm{DUSP}(t)\frac{\mathrm{ERK}^*(t)}{D_5 + \mathrm{ERK}^*}\right)$$

$$\dot{\mathrm{ERK}}^*(t) = \left(k_5 \mathrm{MEK}^*(t)\frac{\mathrm{ERK}(t)}{K_5 + \mathrm{ERK}(t)}\right) - \left(\gamma_5 \mathrm{DUSP}(t)\frac{\mathrm{ERK}^*(t)}{D_5 + \mathrm{ERK}^*}\right)$$

$$\dot{\mathrm{NFB}}(t) = -\left(k_6 \mathrm{ERK}^*(t)\frac{\mathrm{NFB}(t)}{K_6 + \mathrm{NFB}(t)}\frac{\mathrm{EGFR}^*(t)^2}{K_{6R}^2 + \mathrm{EGFR}^*(t)^2}\right) + \left(\gamma_6 \frac{\mathrm{NFB}^*(t)}{D_6 + \mathrm{NFB}^*(t)}\right)$$

$$\dot{\mathrm{NFB}}^*(t) = \left(k_6 \mathrm{ERK}^*(t)\frac{\mathrm{NFB}(t)}{K_6 + \mathrm{NFB}(t)}\frac{\mathrm{EGFR}^*(t)^2}{K_{6R}^2 + \mathrm{EGFR}^*(t)^2}\right) - \left(\gamma_6 \frac{\mathrm{NFB}^*(t)}{D_6 + \mathrm{NFB}^*(t)}\right)$$

$$\dot{\mathrm{dusp}}(t) = \left(\mathrm{dusp}_{basal}\left(1 + \mathrm{dusp}_{ind}\frac{\mathrm{ERK}^{*2}}{K_{dusp} + \mathrm{ERK}^{*2}}\right)\frac{\log_{10}(2)}{\mathrm{T}_{dusp}}\right) - \left(\mathrm{dusp}(t)\frac{\log_{10}(2)}{\mathrm{T}_{dusp}}\right)$$

$$\dot{\mathrm{DUSP}}(t) = \left(\mathrm{dusp}(t)\frac{\log_{10}(2)}{\mathrm{T}_{DUSP}}\right) - \left(\mathrm{DUSP}(t)\frac{\log_{10}(2)}{\mathrm{T}_{DUSP}}\right)$$

TABLE IV. Prior distributions of the Ras-Raf-MEK-ERK parameters.

Parameter	Meaning	Prior interval	$\boldsymbol{\theta}$
k_1	Egf-dependent Egf-Receptor activation rate	[exp(−10), exp(5)]	
γ_1	Egf-Receptor deactivation rate	[exp(−10), exp(5)]	
k_2	Egf-Receptor dependent Ras activation rate	[exp(−10), exp(5)]	
K_2	Hill constant for Ras activation	[exp(−10), exp(5)]	
γ_2	Ras deactivation rate	[exp(−10), exp(5)]	
D_2	Hill constant for Ras deactivation	[exp(−10), exp(5)]	
k_3	Ras dependent Raf activation rate	[exp(−10), exp(5)]	
K_3	Hill constant for Raf activation	[exp(−10), exp(5)]	
K_{NFB}	Coupleing constant for negative feedbock of NFB on Raf activation	[exp(−10), exp(5)]	
γ_3	Raf deactivation rate	[exp(−10), exp(5)]	
D_3	Hill constant for Raf deactivation	[exp(−10), exp(5)]	
k_4	Raf dependent MEK activation rate	[exp(−10), exp(5)]	
K_4	Hill constant for MEK activation	[exp(−10), exp(5)]	
γ_4	MEK deactivation rate	[exp(−10), exp(5)]	
D_4	Hill constant for MEK deactivation	[exp(−10), exp(5)]	
k_5	MEK dependent ERK activation rate	[exp(−10), exp(5)]	
K_5	Hill constant for ERK activation	[exp(−10), exp(5)]	
γ_5	DUSP dependent ERK deactivation rate	[exp(−10), exp(5)]	
D_5	Hill constant for ERK deactivation	[exp(−10), exp(5)]	
k_6	ERK and Egf-Receptor dependent NFB activation rate	[exp(−10), exp(5)]	
K_6	Hill constant for NFB dependent NFB activation	[exp(−10), exp(5)]	
K_{6R}	Hill constant for Egf-Receptor dependent NFB activation	[exp(−10), exp(5)]	
γ_6	NFB deactivation rate	[exp(−10), exp(5)]	
D_6	Hill constant for NFB deactivation	[exp(−10), exp(5)]	
dusp_{basal}	Basal dusp transcription rate	[exp(−10), exp(5)]	
dusp_{basal}	ERK induced dusp transcription rate	[exp(−10), exp(5)]	
K_{dusp}	Hill constant for ERK dependent dusp transcription	[exp(−10), exp(5)]	
T_{dusp}	Half life of dusp	fixed to 90 min	
T_{DUSP}	Half life of DUSP	fixed to 90 min	
$\text{egf}_{input}(t)$	Egf input at time t	input parameter	

2. Noise Model

The measurements are obtained using a FRET sensor (see [27]) and then normalized, such that the average of the signal is 1 at the beginning of the experiment. FRET sensors consist of a sensing unit, that has two fluorescent proteins on it. This sensing unit responds to particular biochemical activity with conformational changes that either change the distance of the two fluorescent proteins of the sensor, or their orientation. In both cases the result is that stimulation with light in the donor channel stimulated the donor protein, which, when excited, partially also stimulates the acceptor protein. Thus upon stimulation in the donor channel one is also able to measure a response in the acceptor channel. This conformational changes in the FRET protein are reversible, such that an increase in the measured acceptor channel upon donor stimulation is proportional to the desired signal to be measured. The observed signal is taken to be the ration of the acceptor response and the donor response upon donor stimulation. For modelling we followed a similar approach as in [3]. We took the final measurement to be

$$R(t) = \lambda \left(\frac{\text{ERK}^*(t) E f_{AA}}{(\text{ERK}(t) + (1-E)\text{ERK}^*(t)) f_{DD}} + \text{bg} \right) + \epsilon_\sigma,$$

where ϵ_σ is a zero mean normal random variable with σ standard deviation. All parameters and their priors can be found in table IX A 2.

3. Inferred Posterior

We ran the LNS algorithm for the Egf-Ras-Raf-Mek-Erk model using trajectories corresponding to 6 differnt input patterns. The chosen trajectories are the means of responses to those different input patterns. The algorithm was run with $\alpha = 50\%$, 50 DP-GMM iterations for the density approximations. The inference was done on the Brutus Cluster of the ETH Zurich (`https://www1.ethz.ch/id/services/list/comp_zentral/cluster/index_EN`). Figure **??** shows

TABLE V. Prior distributions of the FRET Measurement model.

Parameter	Meaning	Prior interval
bg	background fluorescence	[exp(−7), exp(6)]
λ	Scaling constant	[exp(−7), exp(6)]
σ	Standard deviation of the normal noise on the measurement	[exp(−8), exp(−3)]
E	FRET efficiency (energy transferred from stimulated donor molecule to acceptor molecule in active FRET probes)	[exp(−4), exp(−0.5)]
f_{AA}	Fraction of acceptor emission captured by the acceptor channel	[exp(−8), exp(5)]
f_{DD}	Fraction of donor emission captured by the donor channel	[exp(−8), exp(5)]

TABLE VI. Species and intial numbers of the LacGfp model

Species	Notation	Initial Distribution
LacI mRNA	$lacI$	U([0, 5])
LacI protein monomer	$LACI$	U([0, 10])
LacI dimer	$LACI2$	fixed to 0
Unoccupied (active) Lac promoter	$PLac$	fixed to 0
Occupied Lac promoter with 2 repressor molecules bound	$O2Lac$	fixed to 0
Occupied Lac promoter with 4 repressor molecules bound	$O4Lac$	U([50, 70])
GFP mRNA	gfp	fixed to 0
"Dark" GFP protein	GFP	fixed to 0
Mature GFP protein	$mGFP$	fixed to 0

U([a, b]) denotes the uniform distribution on the interval [a, b].

the posteriors obtained for the parameters of the model. Figure **??** shows the posterior for the initial concentrations of the Egf Receptors, the model species NFB and of Dusp mRNA and protein. Figure **??** shows the posterior for the parameters of the FRET model. It can be seen that several parameters are very well identified. Since we don't know the actual value of those parameters, the only way to see if those parameters are reasonable is to look at the resulting simulations (Main paper Figure 3 A and B). The simulations seem to reproduce the behaviour of the system faithfully. To validate if the chosen model is appropriate we could now compare the inferred posterior with measured biological quantities (for instance the initial concentration of Dusp). We can also draw conclusions for the model from the reaction parameters. For instance, the marginal posterior for K_{NFB} is concentrated at around 0.001, which is very low, compared to other parameters. Taking a look at the reaction equations IX A 1, we see that such a low value for K_{NFB} means that the model response to almost any increase in NFB^* in a switch like manor. Similarly, a very large value for D_6 indicates that the deactivation of NFB^* does not seem to play any role. These are just two illustrations what kind of conclusions can be drawn from an available posterior over the parameters and in what way the posterior may aid in further developing an available model.

B. LacGfp example

The second model we use for the demonstration of our algorithm is rather large. It is for a LacGfp system with 9 species and 18 reactions. A schematic representation of the system is given in Figure 7 A. Table IX B shows the species involved and their initial distribution for the simulation.

The reactions of the model all follow mass action kinetics and take the following form:

1. Reactions

1. $\emptyset \xrightarrow{\theta_1} lacI$. Transcription of lacI mRNA (constitutive).

2. $lacI \xrightarrow{\theta_2} \emptyset$. Degradation of lacI mRNA (constitutive).

3. $lacI \xrightarrow{\theta_3} lacI + LACI$. Translation of LACI protein.

TABLE VII. Prior distributions of the Lac-GFP parameters and the real values $\boldsymbol{\theta}^*$ used for the simulation of $\boldsymbol{y}$.

Parameter	Meaning	Prior interval	$\boldsymbol{\theta}^*$
θ_1	lacI transcription rate	[1, 50]	1.215
θ_2	lacI degradation rate	[0.1, 30]	7.430
θ_3	LACI translation rate	[1, 50]	1.439
θ_4	IPTG-independent LACI degradation rate	[0.1, 20]	3.971
θ_5	IPTG-induced increase in LACI degradation rate	[1, 100]	1.173
θ_6	Dimerization rate of LACI	[0.16, 3000]	1716.683
θ_7	Dissociation rate of LACI dimers	[0.1, 20]	6.282
θ_8	Binding rate of LACI dimers to Lac promoter	[0.0016, 1.6]	0.433
θ_9	Dissociation rate of LACI dimers from Lac promoter	[0.1, 20]	0.679
θ_{10}	Tetramerization rate of LACI	[0.16, 1.600]	173.576
θ_{11}	Dissociation rate of LACI tetramers	[0.1, 20]	0.623
θ_{12}	gfp transcription rate from free PLac promoter	[1, 150]	119.232
θ_{13}	gfp transcription rate from LACI dimer-bound PLac	[0.1, 15]	0.120
θ_{14}	gfp transcription rate from LACI tetramer-bound PLac	$[10^{-5}, 1]$	0.006
θ_{15}	gfp degradation rate	[0.1, 20]	1.673
θ_{16}	GFP translation rate	[1, 50]	32.318
θ_{17}	GFP degradation rate	[0.1, 20]	1.196
θ_{18}	GFP maturation rate	[0.1, 20]	2.102

For all the Lac-Gfp inference problems presented in this paper, each parameter was assigned an independent uniform log prior distribution in the interval listed in the table.

4. $LACI \xrightarrow{\theta_u} \emptyset$, where $\theta_u = \theta_4 + \theta_5[IPTG]$. Degradation of LACI protein, increased by the input (IPTG).

5. $LACI + LACI \xrightarrow{\theta_6} LACI2$. Dimerization of LACI protein.

6. $LACI2 \xrightarrow{\theta_7} LACI + LACI$. Dissociation of LACI dimer.

7. $LACI2 + PLac \xrightarrow{\theta_8} O2Lac$. Binding of LACI dimer to Lac operator sequence.

8. $O2Lac \xrightarrow{\theta_9} LACI2 + PLac$. Dissociation of LACI dimer from operator sequence.

9. $O2Lac + O2Lac \xrightarrow{\theta_{10}} O4Lac$. Binding of two LacI/operator complexes and tetramerization.

10. $O4Lac \xrightarrow{\theta_{11}} O2Lac + O2Lac$. Dissociation of tetramer structure.

11. $PLac \xrightarrow{\theta_{12}} PLac + gfp$. Transcription of gfp mRNA from active Lac promoter.

12. $O2Lac \xrightarrow{\theta_{13}} O2Lac + gfp$. Transcription of gfp mRNA from Lac promoter bound to LacI dimer.

13. $O4Lac \xrightarrow{\theta_{14}} O4Lac + gfp$. Transcription of gfp mRNA from Lac promoter bound to LacI tetramer.

14. $gfp \xrightarrow{\theta_{15}} \emptyset$. Degradation of gfp mRNA.

15. $gfp \xrightarrow{\theta_{16}} gfp + GFP$. Translation of dark GFP protein.

16. $GFP \xrightarrow{\theta_{17}} \emptyset$. Degradation of dark GFP protein.

17. $GFP \xrightarrow{\theta_{18}} mGFP$. Maturation of GFP.

18. $mGFP \xrightarrow{\theta_{17}} \emptyset$. Degradation of mature GFP protein.

The effect of IPTG on the system is modelled as an increase in the degradation rate of $LACI$. The relationship between such rate and the inducer concentration, denoted $[IPTG]$, is assumed to be linear. We take the $[IPTG]$ concentration to be 10 μM. There is only $P = 1$ measured species, namely $mGFP$.

The parameters used to simulate the dataset $\boldsymbol{y}$ as well as the priors used for the inference are shown in Table VII.

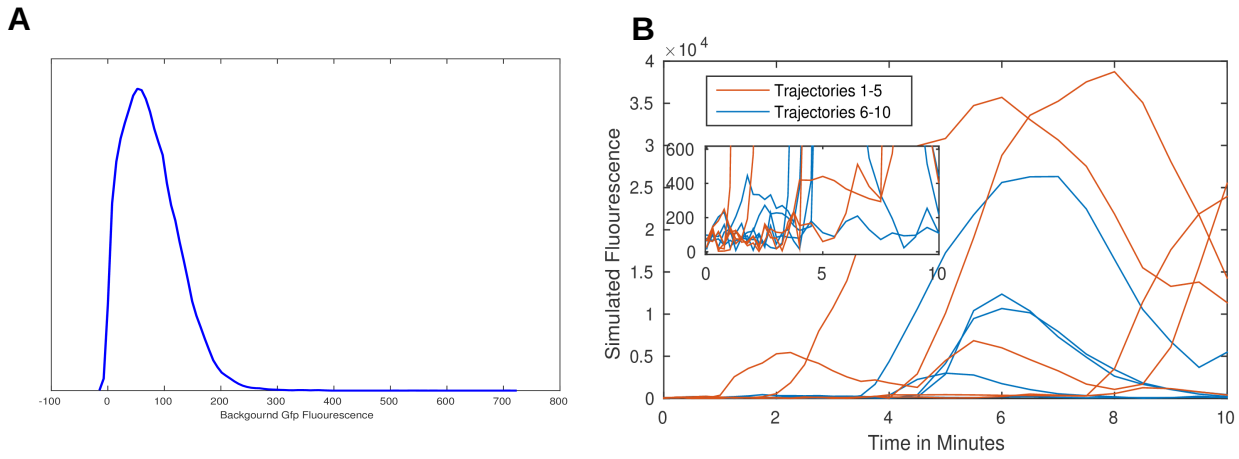

FIG. 8. **A**: The distribution of the background fluorescence. **B**: 10 trajectories of simulated Lac-Gfp data, measured at 29 time points.

We assume that our observation y_τ at a certain time τ is distributed according to

$$y_\tau \sim \mathcal{N}(22.85x_\tau, 5.72\sqrt{x_\tau}) + \mathcal{B},$$

where x_τ is the number of GFP molecules at time τ, $\mathcal{N}(\mu, \sigma)$ is the normal distribution with mean μ and standard deviation σ and $\mathcal{B}$ is a known background fluorescence actually taken from flow cytometry data. The distribution of the background fluorescence is shown in Figure 8 A. This noise model implies a mean fluorescence of 22.85 for each protein and a standard deviation of 5.72. Figure 8 B shows the first 10 simulated trajectories of the Lac-Gfp system measured on 29 timepoints. As can be seen, most trajectories exhibit switch like behaviour. This is indeed one of the properties of the considered example that make it particularly hard to perform likelihood approximation for the single trajectories.

X. TUNING THE ALGORITHM

A. Different values for α

We discuss the choice of α by running the algorithm with different values for the quantile α. The number of particles used for the approximation of the likelihood was 200 and only one observed trajectory was considered. We performed the parameter inference with $\alpha = 10\%, 20\%$ and 50%. The corresponding developments of the minimal and maximal log-likelihood can be seen in Figure 9 B. As expected we can see that choosing α small (10%) results in a slow increase of log-likelihoods, while a higher α corresponds to a steeper increase. The minimal and maximal log-likelihood seems to converge to the same value for each different values of α. Figure Figure 9 A shows the marginals of the posterior obtained using different values of α. We see that, for the most part, each run seems to approximate the same posterior. The minor differences (for instance in the marginal for θ_2 can be explained with a slightly different exploration of the parameter space.

B. Accuracy of likelihood computation

The approximation of the likelihood is what decides whether a particle is accepted or not, thus the accuracy of the likelihood approximation is expected to play an important role for the inference of the posterior. We applied our algorithm to the Lac-Gfp model, using only one observed trajectory and different numbers of particles for the likelihood approximation. The resulting posteriors are shown in Figure 10 A.

As can be seen, for the most part the posteriors do not seem to differ too much. Only for parameter θ_9, the posterior obtained with only 50 particles seems to differ significantly from the one obtained using 200 and 1000

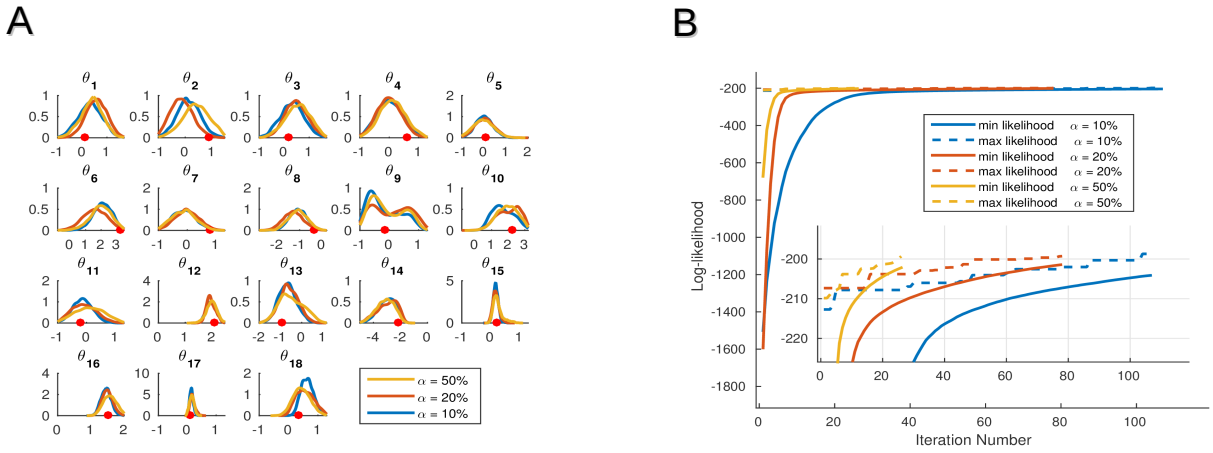

FIG. 9. **A:** Marginal posteriors obtained for the Lac-Gfp model using different values for α. **B:** Minimal and maximal log-likelihood for each level set, for different values for α.

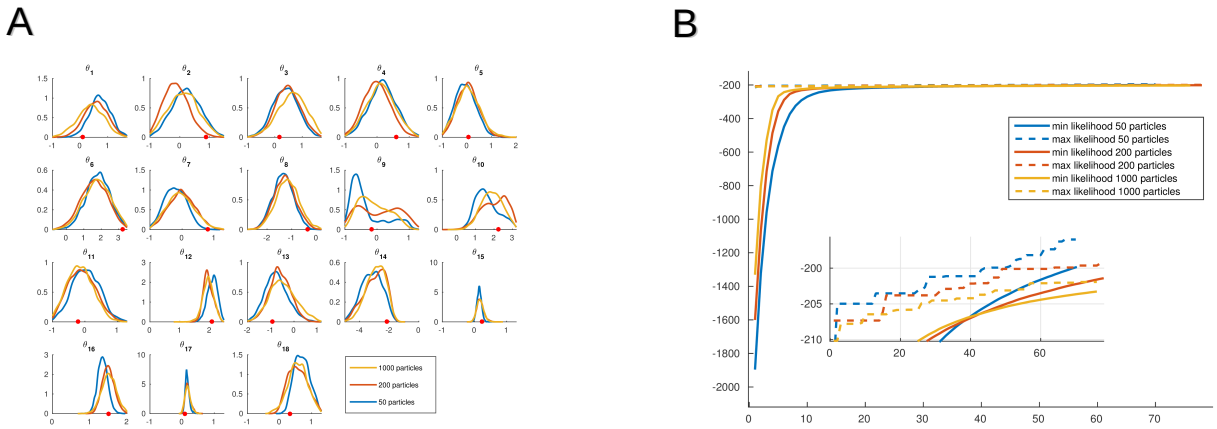

FIG. 10. **A:** Marginal posteriors obtained for the Lac-Gfp model using different number particles for the likelihood approximation. A higher number of particles indicates a more accurate posterior approximation. **B:** Minimal and maximal log-likelihood for each level set, for different particle number for the likelihood approximation.

particles. However, 50 particles are indeed very little and considering the switching behaviour of the system, it is not surprising that the estimated posterior differs from the posterior obtained using 200 or 1000 particles for the likelihood approximation. However, the difference between 200 and 1000 particles for likelihood approximation seems to be very small indicating, that 200 particles are already enough to approximate the posterior accurately. Figure 10 B illustrates the development of the minimal and maximal log-likelihood for each level set for different number of particles for the likelihood approximation. It can be seen that for all particle numbers the gap between minimal and maximal log-likelihood gets smaller. As the likelihood approximation gets more precise (the number of particles increases), the minimal and maximal log-likelihoods tend to get smaller, when compared to less precise approximations.